

Solving Stress Constrained Compliance Minimization using Anisotropic Mesh Adaptation, the Method of Moving Asymptotes and a Global p-norm

Kristian Ejlebjerg Jensen¹

*Imperial College London, Department of Earth Science and Engineering, London SW7 2AZ,
United Kingdom^{a)}*

(Dated: 5 August 2015)

ABSTRACT

The p-norm often used in stress constrained topology optimisation supposedly mimics a delta function and it is thus characterised by a small length scale and ideally one would also prefer to have the solid-void transition occur over a small length scale, since the material in this transition does not have a clear physical interpretation. We propose to resolve these small length scales using anisotropic mesh adaptation. We use the method of moving asymptotes with interpolation of sensitivities, asymptotes and design variables between iterations. We demonstrate this combination for the portal and L-bracket problems with $p=10$, and we are able to investigate mesh dependence. Finally, we suggest relaxing the L-bracket problem statement by introducing a rounded corner.

Keywords: Anisotropic, mesh, adaptation, topology, optimisation, Stress, Constraints, FEniCS, PRAgMaTic.

1. INTRODUCTION

Anisotropic mesh adaptation is an established technique for ensuring computational efficiency in the context of multiscale problems^{1–3}, sometimes reducing the computational work with several orders of magnitude⁴. Within the field of structural optimisation fixed structured meshes remain popular due to ease of implementation and compatibility with mathematical optimisers. It is thus mainly in the context of methods utilising various continuous sensitivities^{5,6} that mesh adaptivity has been applied, and even then mainly in the context of mesh refinement, but some of the most recent advancements also relate to the use of swapping operations⁷. Parallelism is another popular technique for speeding up computations^{8–10} and for truly large scale problems, it indeed is the only option.

The community of structural optimisation has devoted significant effort towards not only finding stiff and light structures, but also structures that do not break under load. That is the structure has to satisfy a stress constraint. Recent strategies include methods with global stress constraints, either with an explicit design representation, such that computation of void elements and related stress is avoided¹¹, or an implicit representation^{12,13}. An essential question in this context is a global versus a local constraint as investigated in the context of level-set methods^{12,13}. One can also choose to let the number of constraints vary throughout the optimisation¹⁴.

An older, but still popular approach^{15,16}, approximates the local stress constraints with a single global constraint using a p-norm. The approximation is good for large values of p , but this causes numerical problems and therefore

a compromise has to be made in practice. If the numerical problems are due to discretisation errors, one can hope to achieve a better compromise with more computational power or by being more efficient with the power available. It is the second option, that we choose to pursue with anisotropic mesh adaptation in this paper. The mesh adaptation introduces inconsistencies, but this drawback can be outweighed by the benefits of improved resolution of small length scales.

We choose to calculate discrete sensitivities as this allows us to harness the power of libraries for automatic adjoint derivation¹⁷. A continuous sensitivity is required for driving the mesh adaptation as well as when interpolating between meshes, but this is trivial to compute by dividing the discrete sensitivity with the design variable volumes, and one can invert the operation, if discrete variables are desired for the optimiser¹⁸.

2. ANISOTROPIC MESH ADAPTATION

Stretched elements with small or large angles are normally discouraged in the context of the finite element method, as they give rise to not only large discretisation errors, but also cause problems with iterative solvers. This wisdom is, however, only applicable to problems without any anisotropic features in the solution. In fact, problems with strong anisotropy require anisotropic meshes for optimal use of the computational degrees of freedom. Note that the problems with large angles and iterative solvers are only severe for extreme anisotropy and, furthermore, it is possible to create anisotropic meshes with few large angles¹⁹.

The properties of the quasi optimal mesh can be derived in a continuous sense³ using a metric tensor, \underline{M} . This is a symmetric and positive definite tensor field that maps the optimal element to the unit element, which has

^{a)}Electronic mail: kristianejlebjerg@gmail.com

unit length edges. The first step is to compute the Hessian, $\underline{\underline{\mathbf{H}}}$, of the variable whose discretisation error is to be minimised and convert it into a positive definite matrix. This is done by taking the absolute value in the principal frame using the operator $\underline{\underline{\text{abs}}}$, which corresponds to removing the sign of the eigenvalues. The optimal mesh metric²⁰ can then be expressed as

$$\underline{\underline{\mathcal{M}}} = \frac{1}{\eta} (\det[\underline{\underline{\text{abs}}}(\underline{\underline{\mathbf{H}}})])^{-\frac{1}{2q+d}} \underline{\underline{\text{abs}}}(\underline{\underline{\mathbf{H}}}), \quad (1)$$

where d is the number of dimensions, q is the error norm to be minimised, \det is the determinant and η is a scaling factor. It is possible to combine the metrics of several variables using the inner ellipse method illustrated in figure 1, see² for implementation details. Note that the unit of η is always so that the metric in equation (1) has dimension of squared inverse length.




FIG. 1. The inner ellipse method is illustrated in the case of intersection, but it is common to see one ellipse entirely within the other, such that anisotropy is preserved.

Once a final metric has been calculated, it can be passed to an anisotropic mesh generator together with the current mesh, in fact our metric is defined on the nodes of the current mesh. We use two mesh generators, which applies four local mesh modifications: Coarsening, refinement, swapping and smoothing, see figure 2. The mesh generators work by applying the modifications to optimise a heuristic quality measure²¹, which quantifies the difference between the discrete mesh and the optimal continuous one. We use an optimised C++ mesh generator, PRAGMaTic²² to illustrate the negligible performance penalty of mesh adaptation, but presently this does not support curved geometries, so in that case we resort to a slower MATLAB/Octave implementation.

Our implementation can generate a metric from a function represented on continuous 2nd order polynomials, but 1st order polynomials can also be used. In the latter case Galerkin projection is used as derivative recovery technique to compute the Hessian. In any case, Galerkin projection is used to convert the element wise constant metric to a node based representation, before it is passed to the mesh generator.

3. TOPOLOGY OPTIMISATION

In an effort to maximise the impact of this paper, we have tried to choose the most popular method for




FIG. 2. Four local mesh modifications are illustrated: Coarsening, refinement, swapping and smoothing. Only the refinement operation is allowed to increase the worst local element quality for the MATLAB/Octave implementation, while the PRAGMaTic mesh generator also allows the worst element quality to deteriorate for the coarsening.

topology optimisation with stress constraints, but we see no reason why other methods could not benefit equally well from the use of mesh adaptation. Similarly other problems with small length scales, such as many convective problems, could most definitely be solved more efficiently using anisotropic mesh adaptation. This being said, we will focus on the density method with SIMP penalisation²³ and a p-norm for relaxing the local stress constraints to a single global constraint¹⁵. The idea is that the local constraint is satisfied in the limit of p going to infinity. In practice a finite value is chosen, but in the following we will show that adaptive meshing allows for $p = 10$ and that this effectively removes kinks in the design. We use a stress penalisation scheme to eliminate problems with void stress²⁴, and the method of moving asymptotes is used to update the design variables²⁵.

A. Mesh Dependence

It is likely that the perfect solution to stress constrained topology optimisation will continue to allude the scientific community, but one might hope that scholars demonstrate their various heuristic schemes in the context of mesh independence such that the uncertainties of objective and constraint functions can be estimated, allowing for techniques to be compared on an objective basis. We have found the selection of papers including quantitative information related to computational cost in the context of stress constrained topology optimisation to be rather scarce^{6,14,26}, and we have been unable to identify any papers containing a discussion of mesh independence. This might be due to it being well known that the stress does not converge for designs with kinks, but in fact it has been shown that the design curvature can be effectively controlled using projection methods²⁷.

It is well known that the p-norm underestimates the maximum stress, but we have deliberately chosen not

to implement a "local stress fix" scheme^{26,28}, that is a scheme that changes the maximum stress to a conservative value such that the local stress constraint is satisfied. This is due to the fact that local stress constraints constitute a much more difficult problem. The alternative is to solve the relaxed stress problem first, and address the local stress problem by employing a conservative maximum stress in combination with a large p-norm. Regardless, the design has to go through a manual post-processing step as it is already typical for designs obtained with topology optimisation, but the topology is normally fixed in the post-processing and therefore it is important that the optimisation result is robust and thus mesh independent.

4. PROBLEM SETUP




FIG. 3. Both the portal (a) and the L-bracket (b) are good benchmarks for stress constrained topology optimisation due to the stress concentration in the corners of the geometries. The load is distributed over the length L_1 in order to avoid a stress singularity in the problem definitions. Sharp corners can also give rise to stress singularities, and we have investigated the effect of this by rounding the corner of the L-bracket. The rollers to the lower right of the portal frame are associated with zero normal displacement, Ω_{ice}

We consider two 2D problems, the L-bracket and the portal frame²⁸, both with finite load and support areas, Ω_{load} and $\Omega_{u=0}$, as illustrated in figure 3. We model plane stress and use a Helmholtz filter²⁹ to compute a filtered design, $\tilde{\rho}$, with a minimum length scale, L_{min} ,

$$\mathbf{0} = \nabla \cdot \underline{\underline{\sigma}}, \quad \underline{\underline{\sigma}} = 2G\underline{\underline{\epsilon}} + \lambda \underline{\underline{I}} (\text{Tr}(\underline{\underline{\epsilon}}) + \partial_z u_z), \quad (2)$$

$$\underline{\underline{\sigma}} = \underline{\underline{\sigma}}_{load} \quad \text{at} \quad \Omega_{load}, \quad \mathbf{u} \cdot \mathbf{n} = 0 \quad \text{at} \quad \Omega_{ice}$$

$$\mathbf{u} = \mathbf{0} \quad \text{at} \quad \Omega_{u=0}$$

$$\underline{\underline{\epsilon}} = \frac{1}{2} (\nabla \mathbf{u} + [\nabla \mathbf{u}]^T), \quad \partial_z u_z = -\frac{\nu}{1-\nu} \nabla \cdot \mathbf{u},$$

$$G = E \frac{1}{2(1+\nu)}, \quad \lambda = E \frac{\nu}{(1+\nu)(1-2\nu)},$$

$$\tilde{\rho} = \rho + \nabla \cdot L_{min}^2 \cdot \nabla \tilde{\rho}, \quad (3)$$

$$E = E_{min} + (E_{max} - E_{min})\rho^{P_E} \quad (4)$$

where \mathbf{u} is the Lagrangian displacement, $\underline{\underline{\sigma}}$ is the stress, $\underline{\underline{\epsilon}}$ is the strain, $\partial_z u_z$ is the out of plane deformation, $\underline{\underline{I}}$ is the identity tensor, Tr is the trace, G is the shear modulus, λ is Lamé's first parameter, E is Young's modulus, ν is the Poisson ratio, ρ is the design variable, and P_E is the SIMP penalisation exponent. Note that we intend to use a sensitivity filter for the compliance. We thus avoid using the filtered design variable in equation (4), it is only calculated for the purpose of driving the mesh adaptation.

The displacement as well as the design variables are discretised with continuous 1st order polynomials, while we use 2nd order when applying the PDE filter (3). The use of discontinuous constant design variables gives rise to excessive stress concentrations in the context of a sensitivity filter, and thus we have found significantly better results (objective functions) using nodal densities. Element wise constant densities might be used with a density filter, but then the issue of dealing with negative filtered design variables in a robust way arises, hence we use a sensitivity filter.

The forward problem defined in equations (2-4) is solved using FEniCS, an open source finite element engine with a high degree of automation³⁰. We use a direct solver for the forward, adjoint and filter problems, and an iterative solver for Galerkin projections between different element types.

We use the current design and displacement to calculate objective and constraint functions,

$$\sigma_{miss} = E_S \sqrt{\epsilon_{xx}^2 + \epsilon_{yy}^2 - \epsilon_{xx}\epsilon_{yy} + 3\epsilon_{12}^2}$$

$$E_S = E_{min} + (E_{max} - E_{min})\rho^{P_S} \quad (5)$$

$$V = \int_{\Omega} \rho d\Omega$$

$$C = \frac{1}{C_{max}} \int_{\partial\Omega_{load}} \mathbf{u} \cdot \underline{\underline{\sigma}}_{load} \cdot \hat{\mathbf{n}} ds - 1$$

$$S = \left[\int_{\Omega} (\sigma_{miss}/\sigma_{max})^p \right]^{1/p} - 1,$$

where C_{max} is the maximum compliance, P_S is the stress penalisation exponent, σ_{miss} is the von mises stress and

σ_{\max} is its maximum value. The interpolations of the Young's modulus in equations (4) and (5) agree in the limits of $\rho = 0$ and $\rho = 1$ corresponding to void and material, respectively, but one has to penalise the intermediate values differently for stress constrained optimisations²⁴. The discrete gradient of V can be calculated explicitly, and we use dolfin-adjoint¹⁷ to calculate the gradients of S and C . These two discrete gradients are converted to continuous ones by division with the gradient of V and anisotropic Helmholtz smoothing is applied to the stress sensitivity. This smoothing is defined on the continuous level using an equation identical to (3) with a tensor version of L_{\min} , which is based on the Steiner ellipse of the elements. We use the Helmholtz filter (3) for the compliance sensitivity, that is we do not use a sensitivity filter related to non-local elasticity as discussed in³¹. These continuous and smooth sensitivities are then used to calculate metrics associated with the compliance and stress constraints. The filtered design variable is used to calculate a metric associated with the volume constraint and all three metrics are combined using the inner ellipse method, see figure 1. After the metric has been used to adapt the mesh, the optimiser variables (the asymptotes, continuous sensitivities, previous, and previous previous design variables) are interpolated on to the new mesh. Due to extrapolation at curved boundaries it can become necessary to enforce the box constraints on the asymptotes and design variables after this interpolation step. Once this interpolation step is complete, the continuous sensitivities are converted to discrete ones, and the optimiser is called to update the design variables. This completes the optimisation loop as illustrated in figure 4. This methodology is inconsistent from a mathematical programming point of view, but except for the compliance sensitivity filter, we expect the errors to decrease with mesh refinement.

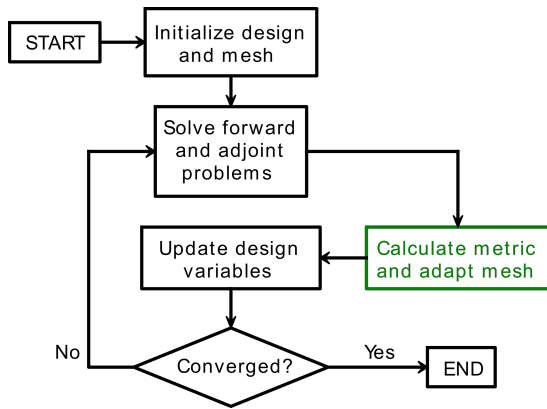


FIG. 4. This flowchart shows the position of the mesh adaptation in the design optimisation procedure, between the adjoint problem and the optimiser.

We consider volume minimisation under stress and compliance constraints, the formal problem statement

being

$$\min_{0 \leq \rho \leq 1} V, \quad s.t. \\ C \leq 0, \quad S \leq 0 \quad \text{and} \quad \text{eq. (2-4)}$$

The optimisation is initialised with a completely solid design, $\rho_0 = 1$.

We make the problem dimensionless using L_{char} and E_{max} as characteristic length scale and stress, respectively. This is reflected in the set of parameters used in the optimisations,

$$L_1 = 0.1L_{\text{char}}, \quad \underline{\sigma}_{\text{load}} = E_{\text{max}}/L_{\text{char}}, \quad \nu = 0.3, \\ L_{\min} = 5 \cdot 10^{-2}L_{\text{char}}, \quad \sigma_{\max} = 1.5E_{\text{max}}, \quad p = 10, \\ P_E = 3, \quad P_S = 0.5, \quad \text{and} \quad E_{\min} = 10^{-3}E_{\text{max}}.$$

The maximum compliance and two of the length scales differ between the two benchmarks as shown in table I. Ideally the values for the maximum stress and compliance should be chosen according to the application, but for the academic benchmarks considered here, we just choose a maximum compliance that results in a non-extreme volume fraction and a maximum stress that results in the stress constraint being active.

parameter \ geometry	portal	L-bracket
L_x/L_{char}	2	1.5
L_y/L_{char}	1	1.5
$C_{\max}L_{\text{char}}^{-2}/E_{\max}$	0.75	2.5
$L_{2,x}/L_{\text{char}}$	2-1/6	
$L_{2,y}/L_{\text{char}}$	0.6	

TABLE I. The L_x and L_y length scale differs between the portal and L-bracket problems, and the maximum compliance is significantly higher for the L-bracket problem.

In addition to these we have the scaling factor, η , and four numerical parameters,

$$c_{\text{MMA}} = 10^3, \quad q = 2, \quad h_{\min} = 10^{-3}, \quad \text{and} \\ \text{AR}_{\max} = 50,$$

where the c_{MMA} is related to enforcement of constraints, q is the error norm to be minimised, h_{\min} is the minimum element edge length and AR_{\max} is the maximum element aspect ratio. Finally, we use move limits, $\Delta\rho$, for the optimiser, which enforce the constraints

$$\text{abs}(\rho_{i+1} - \rho_i) \leq \Delta\rho,$$

and we fix these at $\Delta\rho = 0.1$.

5. RESULTS

In order to investigate mesh independence, we choose the scaling factor related to the metric of the compliance sensitivity as the primary numerical parameter to be varied and scale the number of iteration it_{\max} , the

move limits and the other scaling factors with a dimensionless version of this, $\eta_{\bar{\rho}}$,

$$\begin{aligned} it_{\max} &= \text{round} \left(600 \sqrt{0.015/\eta_{\bar{\rho}}} \right) \\ \eta_C &= \eta_{\bar{\rho}} \\ \eta_S &= 4\eta_{\bar{\rho}}, \end{aligned} \quad (6)$$

where the **round** function rounds the input to the nearest integer. Note the factor of four in equation (6), which reflects a lower emphasis on resolving the sensitivity of the stress constraint. It is possible to optimise with $\eta_S = \eta_{\bar{\rho}}$, but this increases the tendency of the method to produce mesh dependent topologies. This observation and the factor of four is probably sensitive to the choice of $p = 10$.

The sensitivity filter and interpolation of internal optimiser variables between meshes introduces inconsistencies in our implementation, and therefore we do not expect convergence in a strict sense, so we just plot the design variables and mesh elements for the iteration corresponding to the smallest volume at which the constraints are satisfied to the tolerance of the MMA c parameter. The design topology in the final iteration is, however, identical to the one shown³³.

For reference we have performed optimisations with a large maximum stress ($\sigma_{\max} = 15E_{\max}$) to mimic the results of pure compliance minimisation. This is shown in figure 5, which shows the design kinks at the concave corners that we wish to eliminate by imposing the stress constraint.

For the portal problem, we have performed optimisations with $\eta_{\bar{\rho}} = 0.03, 0.015$ and 0.0075 as shown in figure 6. With continuous and linear displacements we get element wise constant strains, so we have to use Galerkin projection to compute continuous and linear stresses, as shown in 6(d-f). The finer optimisations most likely find local asymmetric minima, but in all cases the design kink is eliminated.

For the L-bracket problem, we have performed an optimisation with a fixed mesh and $it_{\max} = 600$, as shown in figure 7. This is to serve as a reference for the optimisations based on adaptive meshing. These are performed with a sharp ($r = 0$) and a rounded corner ($r = 0.01L_{\text{char}}$) for $\eta_{\bar{\rho}}$ equal to $0.03, 0.015$ and 0.0075 as shown in figure 8. The two coarser optimisations with a sharp corner have a component in compression at the load, which might be unstable to perturbations in the load direction. This behaviour has also been observed previously^{16,26}, and most likely it can be fixed by using a second load case. It takes more iterations to find the best feasible design for the sharp corner, possibly indicating that the optimiser has an easier time dealing with the rounded corner, which also might explain the mesh dependent designs for the sharp corner.

It is well known that structures become weaker as the mesh is refined, i.e. the compliance converges from below²⁹ and the same is true for the stress. One would thus expect the volume to converge from below in a stress

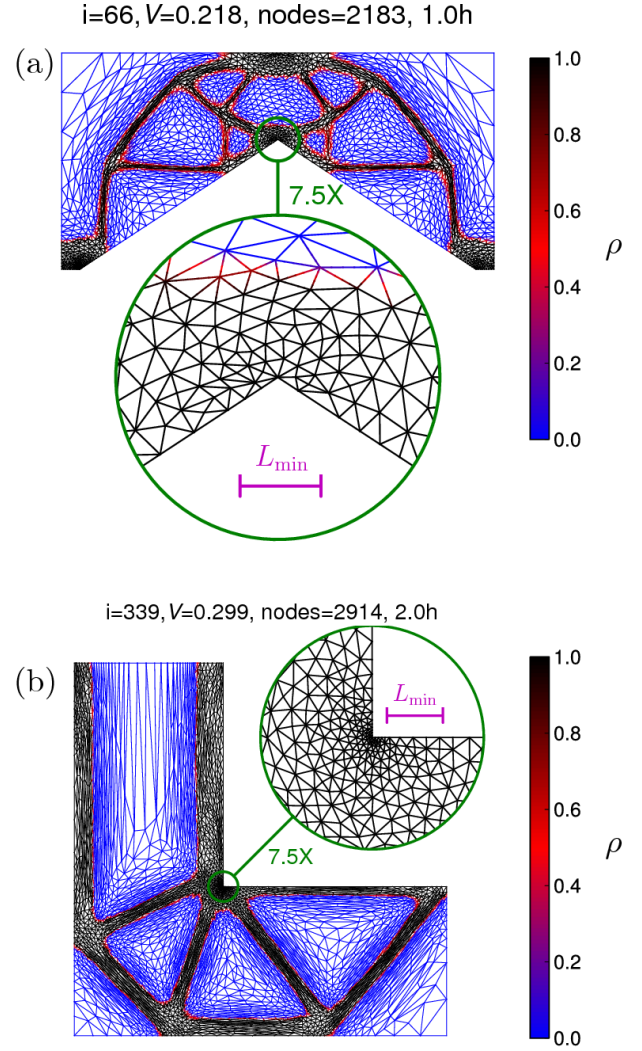


FIG. 5. The result of optimisations using $\sigma_{\max} = 15E_{\max}$ and $\eta_{\bar{\rho}} = 0.03$ are plotted. The large maximum stress results in designs similar to compliance minimisation, thus the bar going horizontally from the load in the L-bracket.

and compliance constrained optimisation problem, but this is not what we see in figure 9, where the objective function is plotted throughout the optimisations for the sharp as well as the rounded corner. We attribute this to the sensitivity filter and the continuous design variables, as this combination causes the area of intermediate design variables to decrease as the mesh is refined. Note that the coarse optimisation with a rounded corner seems to become infeasible and never recover, which might be attributed to numerical inconsistencies related to the sensitivity filter.

The stresses for the design in figure 8 are plotted in figure 10. The maximum stress is not well resolved for $\eta_{\bar{\rho}} = 0.03$, but it does not change much between $\eta_{\bar{\rho}} = 0.015$ and $\eta_{\bar{\rho}} = 0.0075$. It seems like a safety margin no smaller than 2.5 is required, and as such figure 10 highlights the need for post processing and verification

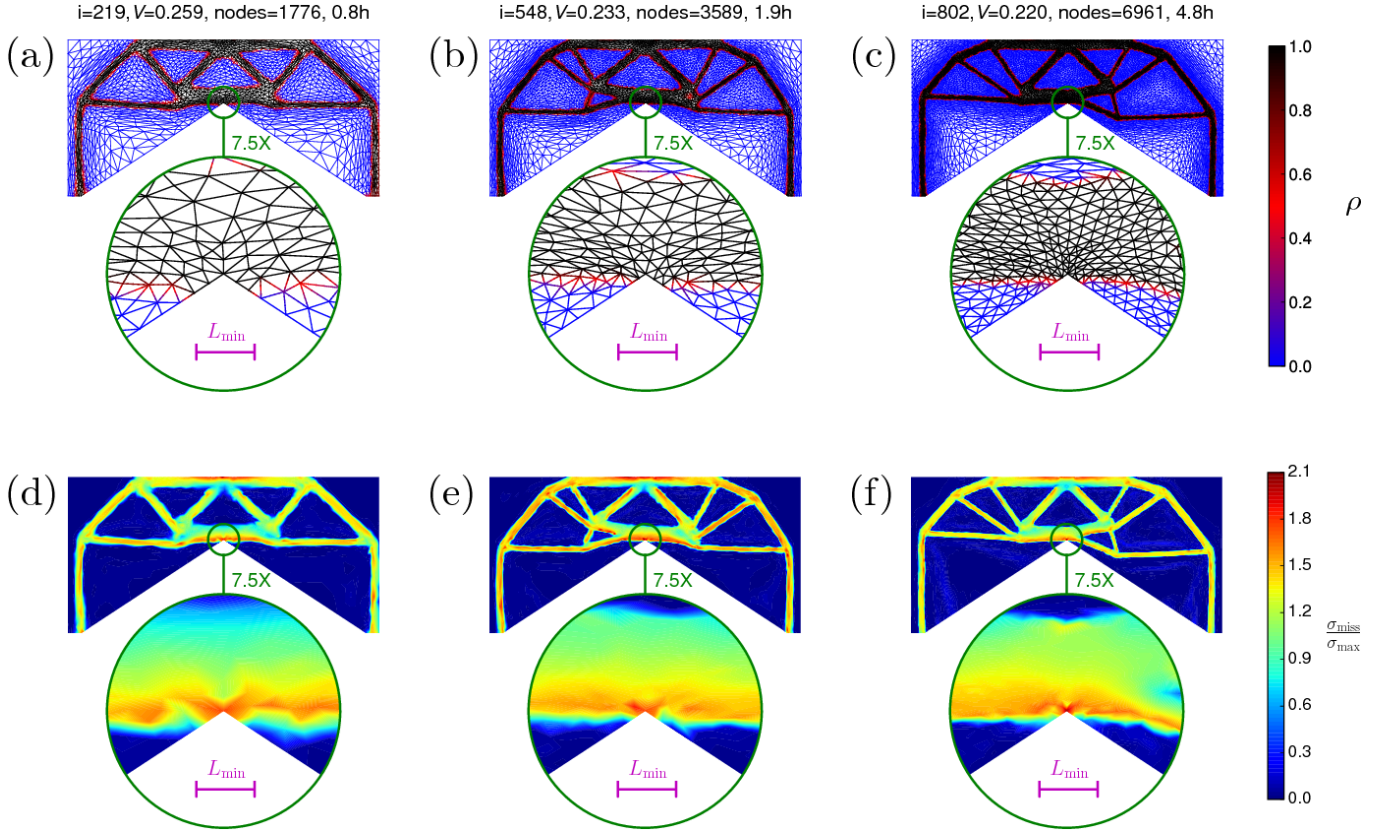


FIG. 6. The portal problem is optimised with $\eta_{\bar{\rho}} = 0.03$ (a,d), 0.015 (b,e) and 0.0075 (c,f). Note the local minimum for the finer optimisations, which is indicative of local minima.

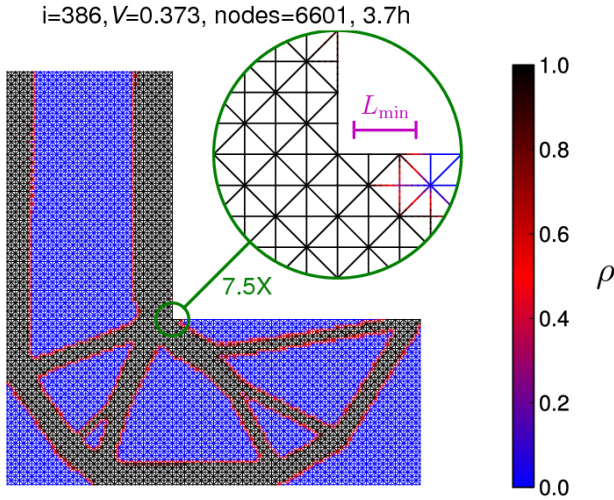


FIG. 7. The result of stress minimisation on a fixed mesh is shown for comparison with the results based on adaptive meshing.

of designs obtained using stress constrained topology optimisation.

Finally, the results of a preliminary study with an iterative solver is shown in table II. The study indicates

that the large element angles do not prohibit the use of iterative solvers, which is important for parallelism and efficiency in three dimensions.

solver \ $\eta_{\bar{\rho}}$	0.03	0.015	0.0075	NA
Direct (LU)	0.4 s	0.72 s	1.47	0.49 s
Iterative (CG+AMG)	0.61 s	1.27 s	2.69	0.68 s

TABLE II. The computational time for solving the forward problem using a direct solver (LU) for the designs with a rounded corner in figure 8. The computational time for solving the same linear systems with the conjugate gradient method and algebraic multigrid as preconditioner is also tabulated. We also show the same timings for the static mesh ($\eta_{\bar{\rho}} = \text{NA}$) in figure 7. This demonstrates that the large element angles do not cause the relative performance of iterative solvers to decrease noticeably.

Computational cost All computations are single threaded, but it is possible to perform mesh adaptation in parallel, see^{22,32}.

We terminate the optimisations using an iteration limit only, and the total computation time[?] before this triggers is shown in hours at the top of each plot in figure 8. We use a direct solver for the filter operations as well as for the forward and adjoint problems. The coarse optimisation of the L-bracket with a rounded corner probably only took around an hour to find the best design, but

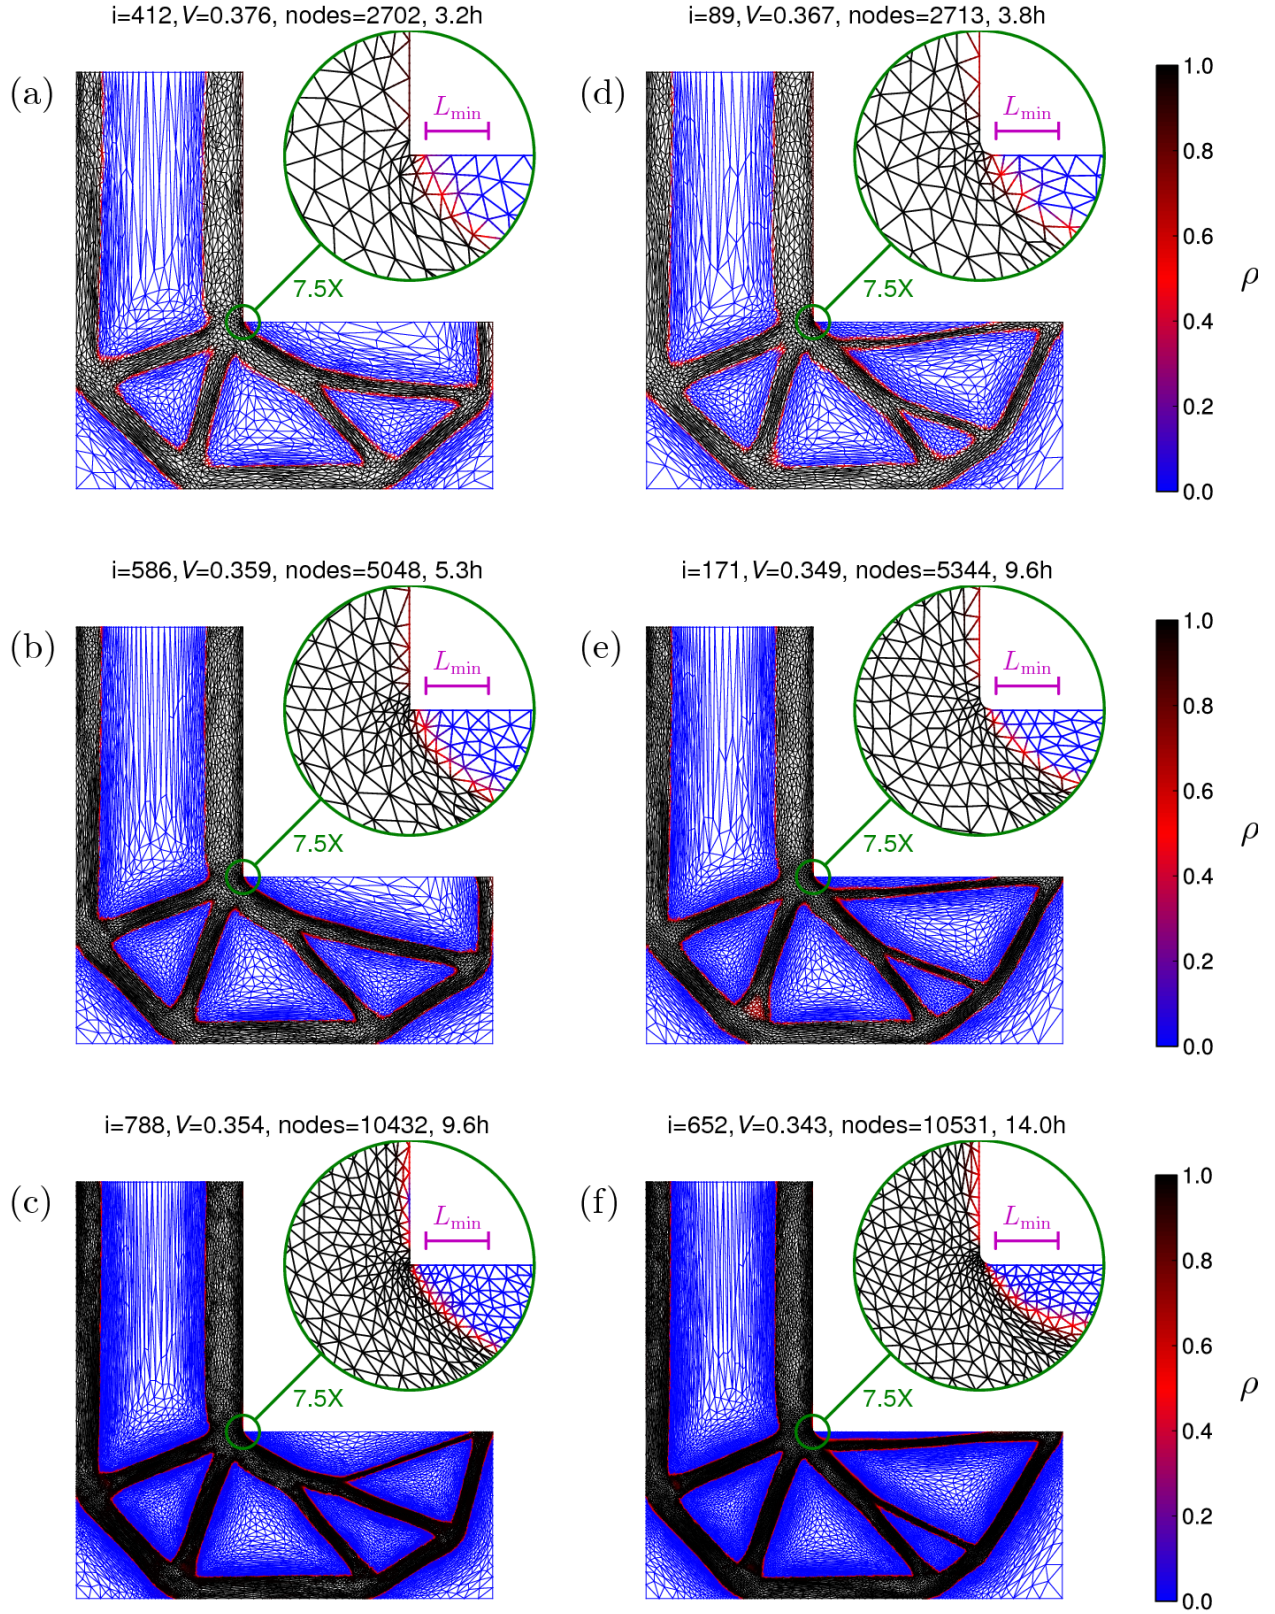


FIG. 8. Optimisations with a sharp (left) and a rounded corner (right) for η_p equal to 0.03 (top), 0.015 (middle) and 0.0075 (bottom). The design variables and mesh elements are shown for the iteration (i) at which the lowest volume fraction (V) occurs, while the stress and compliance constraints are satisfied. See online version for colours.

others have found similar designs with an order of magnitude less effort^{6,14} and comparing figures 7 and 8(a-c)

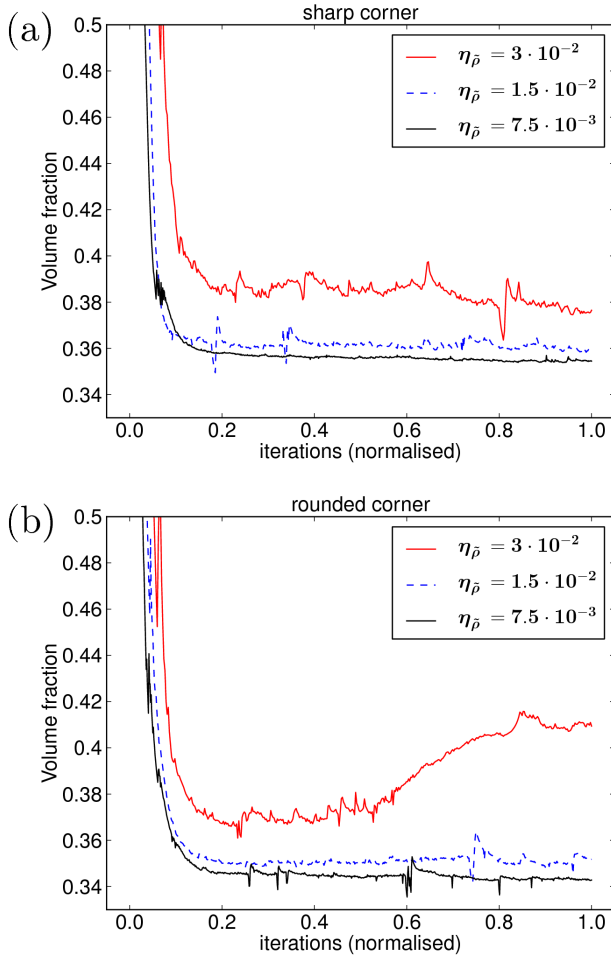


FIG. 9. The volume fraction is plotted versus normalised iteration numbers for different values of $\eta_{\bar{\rho}}$ in the case of a sharp (a) as well as a rounded corner (b). In all cases there are some oscillations, which are probably caused by infeasible designs.

reveals that it is possible to achieve a similar volume fraction using a fixed mesh and a similar amount of computational time. There is thus reason to believe that the presented method needs improvement, if it is to be competitive. It is, however, worth noting that the resolution of the design near the corner is significantly worse for the static mesh, even compared to the most coarse adaptive optimisation, so the fixed mesh might be benefiting from under resolving the stress singularity.

The optimisations with sharp corners are performed with PRaGMaTic as mesh adaptation library, which is an optimised C++ implementation, so the actual mesh adaptation only takes 1-2% of the total computation time. The optimisations with a rounded corner use an almost identical Octave/MATLAB implementation, which (although fully vectorised) is significantly slower, and therefore the mesh generation takes 20-30% of the total computation time. Note that although the meshes conform well to the designs in figure 8, this is not true

for the intermediate iterations. This is due to the fact that we optimise with $p = 10$, so small imperfections in the designs can cause concentrations of stress away from the corner, which lead to many small elements in areas that are close to becoming void.

We calculate the Steiner ellipse for each element, and use this to calculate the element aspect ratio by dividing the radii product with the square of the smaller radius. This element aspect ratio is between 4 and 5 on average throughout the optimisations, which is an indication of speedup due to the anisotropic elements. One can expect the square of this in 3D corresponding to an order of magnitude reduction in computational cost compared to isotropic adaptation.

6. CONCLUSION

We have performed stress constrained topology optimisations using a combination of anisotropic mesh adaptation and the method of moving asymptotes with interpolation of the asymptotes between iterations. We find that it is necessary to use continuous linear design variables and a sensitivity filter. The computational cost of introducing the mesh adaptation is negligible for an optimised implementation.

We argue that it might be beneficial to relax the problem statement to a rounded corner, as a radius of just 1% of the characteristic length scale helps the optimiser find better designs. At least, in this case we are able to demonstrate mesh independence.

7. SUGGESTIONS FOR FUTURE WORK

It would be interesting to see the method applied to design of compliant mechanisms or meta materials due to the fact that the stress concentration results from the nature of the problem rather than the geometry.

Another interesting aspect is the post-processing of the result, i.e. to which extent is it beneficial to have many small elements aligned with the design contour, when the design is to be extracted?

Furthermore, the fact that anisotropic elements can give rise to large angles which cause problems for iterative solvers should be quantified in three dimensions and possibly also addressed by means of a meshing technique relying on advancing fronts¹⁹. Another point for three dimensions is, that due to the trend of increasing number of cores in workstations, the fastest methods generally rely on shared memory parallelism⁷. Therefore one ought to investigate the benefit of parallelism, when extending the method to three dimensions.

Finally, the meshes that occur during the optimisation loop have only small variations during the later stages of the procedure, meaning that the degrees of freedom are chosen efficiently in the spatial dimension only, not in the "optimisation dimension". We suggest to address this

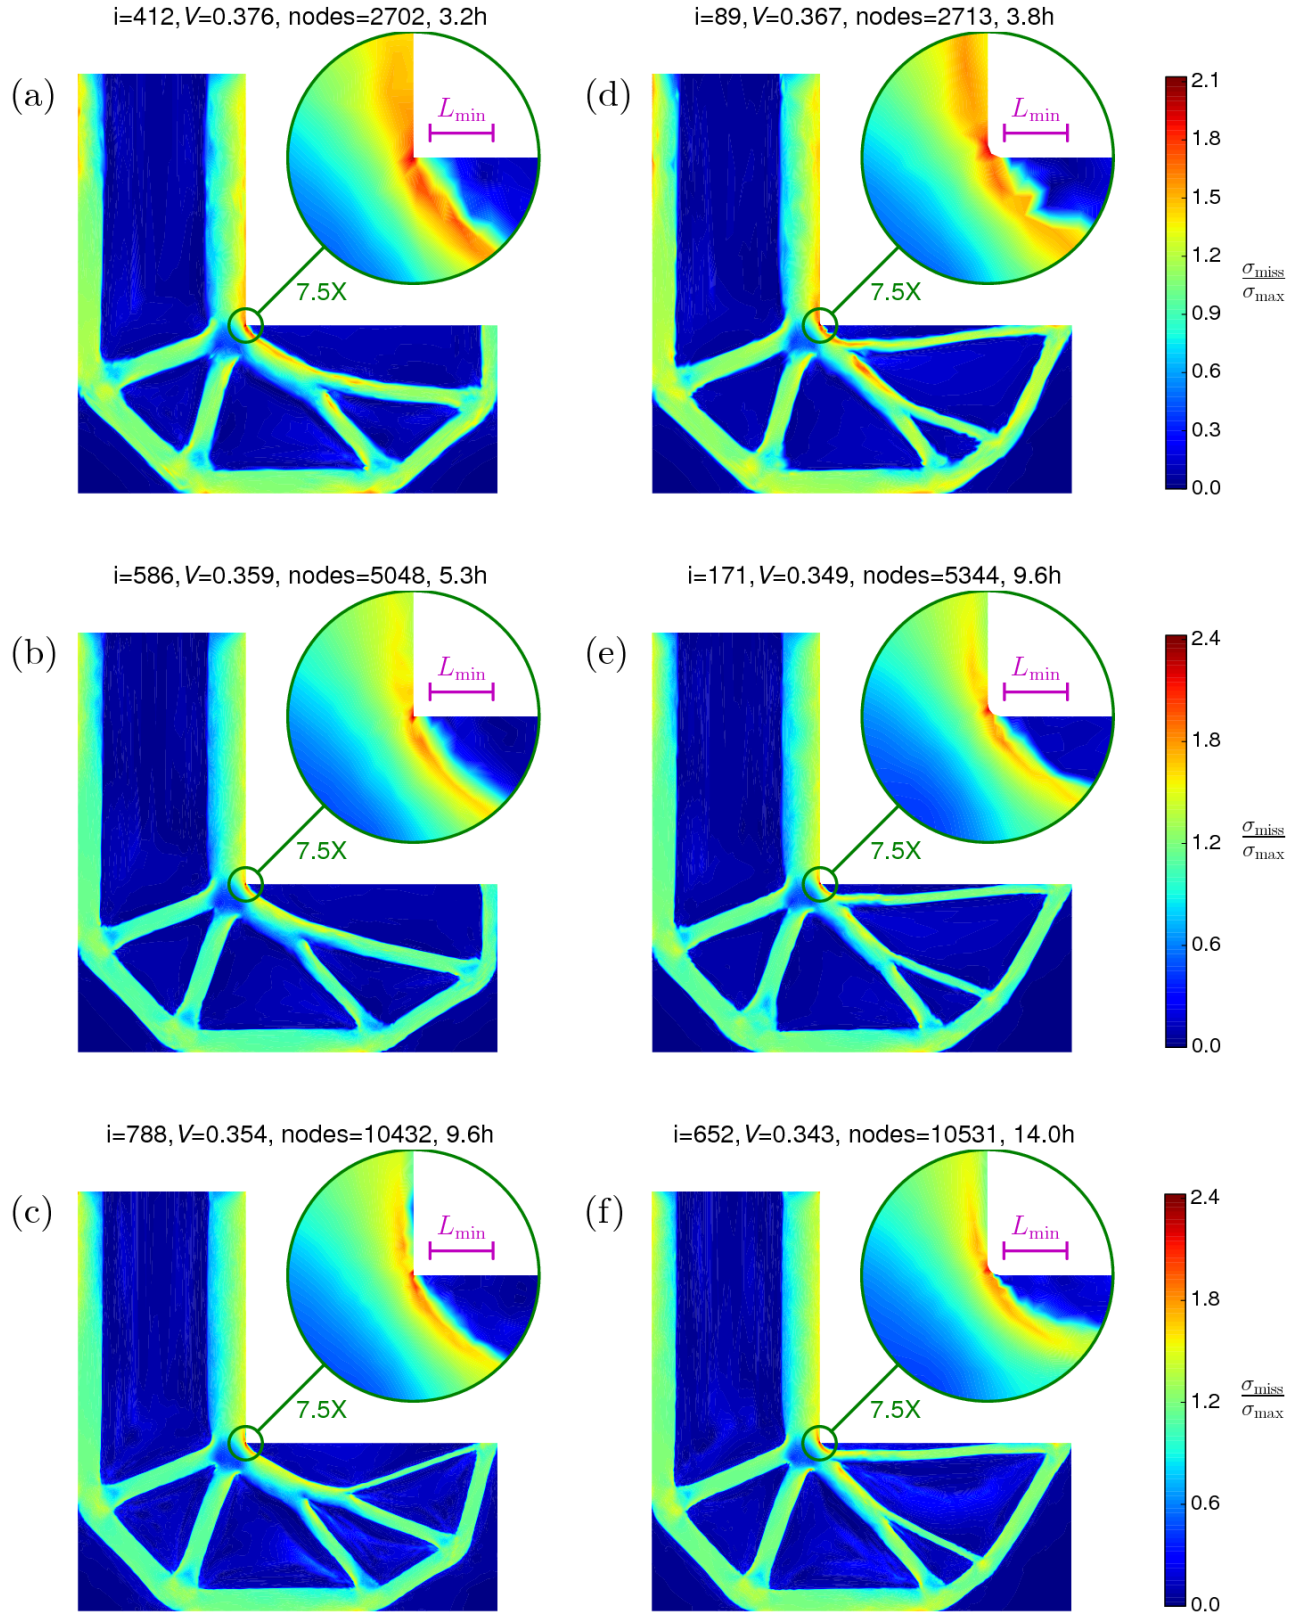


FIG. 10. Optimisations with a sharp (left) and a rounded corner (right) for $\eta_{\bar{p}}$ equal to 0.03 (top), 0.015 (middle) and 0.0075 (bottom). The stress is shown for the iteration (i) at which the lowest volume fraction (V) occurs, while the stress and compliance constraints are satisfied. The stress calculated on the nodes, and the maximum values, in units of the maximum stress, are 2.11, 2.41 and 2.43 for the sharp corner and 2.05, 2.40 and 2.30 for the rounded corner. See online version for colours.

issue using time-space elements and an optimiser defined at the continuous level.

8. ACKNOWLEDGEMENT

This work is supported by the Villum Foundation.

- ¹Wagdi G Habashi, Julien Dompierre, Yves Bourgault, Djaffar Ait-Ali-Yahia, Michel Fortin, and Marie-Gabrielle Vallet. Anisotropic mesh adaptation: Towards user-independent, mesh-independent and solver-independent cfd. part i: General principles. *International Journal for Numerical Methods in Fluids*, 32(6):725–744, 2000.
- ²CC Pain, AP Umpleby, CRE De Oliveira, and AJH Goddard. Tetrahedral mesh optimisation and adaptivity for steady-state and transient finite element calculations. *Computer Methods in Applied Mechanics and Engineering*, 190(29):3771–3796, 2001.
- ³Adrien Loseille and Frédéric Alauzet. Continuous mesh framework part i: well-posed continuous interpolation error. *SIAM Journal on Numerical Analysis*, 49(1):38–60, 2011.
- ⁴Adrien Loseille, Alain Dervieux, and Frédéric Alauzet. Fully anisotropic goal-oriented mesh adaptation for 3d steady euler equations. *Journal of computational physics*, 229(8):2866–2897, 2010.
- ⁵Mathias Wallin, Matti Ristinmaa, and Henrik Askfelt. Optimal topologies derived from a phase-field method. *Structural and Multidisciplinary Optimization*, 45(2):171–183, 2012.
- ⁶Samuel Amstutz and Antonio A Novotny. Topological optimization of structures subject to von mises stress constraints. *Structural and Multidisciplinary Optimization*, 41(3):407–420, 2010.
- ⁷Asger Nyman Christiansen, J Andreas Bærentzen, Morten Nobel-Jørgensen, Niels Aage, and Ole Sigmund. Combined shape and topology optimization of 3d structures. *Computers & Graphics*, 46:25–35, 2015.
- ⁸Thomas Borrvall and Joakim Petersson. Large-scale topology optimization in 3d using parallel computing. *Computer methods in applied mechanics and engineering*, 190(46):6201–6229, 2001.
- ⁹Niels Aage, Erik Andreassen, and Boyan Stefanov Lazarov. Topology optimization using petsc. *Structural and Multidisciplinary Optimization*, 2014.
- ¹⁰Niels Aage, Thomas H Poulsen, Allan Gersborg-Hansen, and Ole Sigmund. Topology optimization of large scale stokes flow problems. *Structural and Multidisciplinary Optimization*, 35(2):175–180, 2008.
- ¹¹Qi Xia, Tielin Shi, Shiyuan Liu, and Michael Yu Wang. A level set solution to the stress-based structural shape and topology optimization. *Computers & Structures*, 90:55–64, 2012.
- ¹²Xu Guo, Wei Sheng Zhang, Michael Yu Wang, and Peng Wei. Stress-related topology optimization via level set approach. *Computer Methods in Applied Mechanics and Engineering*, 200(47):3439–3452, 2011.
- ¹³Wei Sheng Zhang, Xu Guo, Michael Yu Wang, and Peng Wei. Optimal topology design of continuum structures with stress concentration alleviation via level set method. *International Journal for Numerical Methods in Engineering*, 93(9):942–959, 2013.
- ¹⁴Matteo Bruggi and Pierre Duysinx. Topology optimization for minimum weight with compliance and stress constraints. *Structural and Multidisciplinary Optimization*, 46(3):369–384, 2012.
- ¹⁵Pierre Duysinx and Ole Sigmund. New developments in handling stress constraints in optimal material distribution. In *Proc of the 7th AIAA/USAF/NASA/ISSMO Symp on Multidisciplinary Analysis and Optimization*, volume 1, pages 1501–1509, 1998.
- ¹⁶Erik Holmberg, Bo Torstenfelt, and Anders Klarbring. Stress constrained topology optimization. *Structural and Multidisciplinary Optimization*, 48(1):33–47, 2013.
- ¹⁷Patrick E Farrell, David A Ham, Simon W Funke, and Marie E Rognes. Automated derivation of the adjoint of high-level transient finite element programs. *SIAM Journal on Scientific Computing*, 35(4):C369–C393, 2013.
- ¹⁸Preliminary results related to this work was presented at FEniCS 14 in Paris, June 2014 and at the International Meshing Roundtable 23 in London, October 2014.
- ¹⁹Adrien Loseille. Metric-orthogonal anisotropic mesh generation. *Procedia Engineering*, 82:403–415, 2014.
- ²⁰Long Chen, Pengtao Sun, and Jinchao Xu. Optimal anisotropic meshes for minimizing interpolation errors in $^{\wedge}$ -norm. *Mathematics of Computation*, 76(257):179–204, 2007.
- ²¹Yu V Vasilevski and KN Lipnikov. Error bounds for controllable adaptive algorithms based on a hessian recovery. *Computational Mathematics and Mathematical Physics*, 45(8):1374–1384, 2005.
- ²²Georgios Rokos, Gerard J Gorman, James Southern, and Paul HJ Kelly. A thread-parallel algorithm for anisotropic mesh adaptation. *arXiv preprint arXiv:1308.2480*, 2013.
- ²³Martin Philip Bendsoe and Ole Sigmund. *Topology optimization: theory, methods and applications*. Springer, 2003.
- ²⁴GD Cheng and Xiao Guo. ϵ -relaxed approach in structural topology optimization. *Structural Optimization*, 13(4):258–266, 1997.
- ²⁵Krister Svanberg. The method of moving asymptotes a new method for structural optimization. *International journal for numerical methods in engineering*, 24(2):359–373, 1987.
- ²⁶Krishnan Suresh and Meisam Takalloozadeh. Stress-constrained topology optimization: a topological level-set approach. *Structural and Multidisciplinary Optimization*, 48(2):295–309, 2013.
- ²⁷Fengwen Wang, Boyan Stefanov Lazarov, and Ole Sigmund. On projection methods, convergence and robust formulations in topology optimization. *Structural and Multidisciplinary Optimization*, 43(6):767–784, 2011.
- ²⁸Chau Le, Julian Norato, Tyler Bruns, Christopher Ha, and Daniel Tortorelli. Stress-based topology optimization for continua. *Structural and Multidisciplinary Optimization*, 41(4):605–620, 2010.
- ²⁹Boyan Stefanov Lazarov and Ole Sigmund. Filters in topology optimization based on helmholtz-type differential equations. *International Journal for Numerical Methods in Engineering*, 86(6):765–781, 2011.
- ³⁰Anders Logg, Kent-Andre Mardal, Garth N. Wells, et al. *Automated Solution of Differential Equations by the Finite Element Method*. Springer, 2012.
- ³¹Ole Sigmund and Kurt Maute. Sensitivity filtering from a continuum mechanics perspective. *Structural and Multidisciplinary Optimization*, 46(4):471–475, 2012.
- ³²Patrick E Farrell. *Galerkin projection of discrete fields via super-mesh construction*. PhD thesis, Imperial College London, 2009.
- ³³Using an Intel(R) Core(TM) i7 870 @ 2.93GHz.